



Intel[®] Celeron[®] D Processors 335, 330, 325, and 320^Δ

Specification Update

— On 90 nm Process and in the 478-Pin Package

June 2004

Notice: The Intel[®] Celeron[®] D Processor on 90 nm process and in the 478-pin package may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

Document Number: 302354-001



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Celeron® D Processor on 90 nm Process and in the 478-Pin Package may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

⁴Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Over time processor numbers will increment based on changes in clock, speed, cache, FSB, or other features, and increments are not intended to represent proportional or quantitative increases in any particular feature. Current roadmap processor number progression is not necessarily representative of future roadmaps. See www.intel.com/products/processor_number for details.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Intel, Pentium, Intel Xeon, Intel NetBurst, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2004, Intel Corporation



Contents

Revision History	4
Preface	5
Errata.....	13
Specification Changes	27
Specification Clarifications	29
Documentation Changes	31

Revision History

Revision Number	Description	Date of Revision
-001	• Initial Release	June 2004

§

Preface

This document is an update to the specifications contained in the documents listed in the following Affected Documents/Related Documents table. It is a compilation of device and document errata and specification clarifications and changes, and is intended for hardware system manufacturers and for software developers of applications, operating system, and tools.

Information types defined in the Nomenclature section of this document are consolidated into this update document and are no longer published in other documents. This document may also contain information that has not been previously published.

Affected Documents

Document Title	Document Number
<i>Intel® Celeron® D Processor on 90 nm Process and in the 478-Pin Package Datasheet</i>	302253-001

Related Documents

Document Title	Document Location
<i>IA-32 Intel® Architecture Software Developer's Manual Volume 1: Basic Architecture</i>	http://developer.intel.com/design/pentium4/manuals/253665.htm
<i>IA-32 Intel® Architecture Software Developer's Manual Volume 2A: Instruction Set Reference Manual A–M</i>	http://developer.intel.com/design/pentium4/manuals/253666.htm
<i>IA-32 Intel® Architecture Software Developer's Manual Volume 2B: Instruction Set Reference Manual, N–Z</i>	http://developer.intel.com/design/pentium4/manuals/253667.htm
<i>IA-32 Intel® Architecture Software Developer's Manual Volume 3: System Programming Guide</i>	http://developer.intel.com/design/pentium4/manuals/253668.htm

Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics (e.g., core speed, L2 cache size, package type, etc.) as described in the processor identification information table. Care should be taken to read all notes associated with each S-Spec number

Errata are design defects or errors. Errata may cause the Intel® Celeron® D processor on 90 nm process and in the 478-pin package behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

Specification Changes are modifications to the current published specifications. These changes will be incorporated in the next release of the specifications.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specifications.

Documentation Changes include typos, errors, or omissions from the current published specifications. These changes will be incorporated in the next release of the specifications.

General Information

This section contains top marking information for the Intel® Celeron® D processor on 90 nm process and in the 478-pin package.

Figure 1. Example of Sample Package Markings

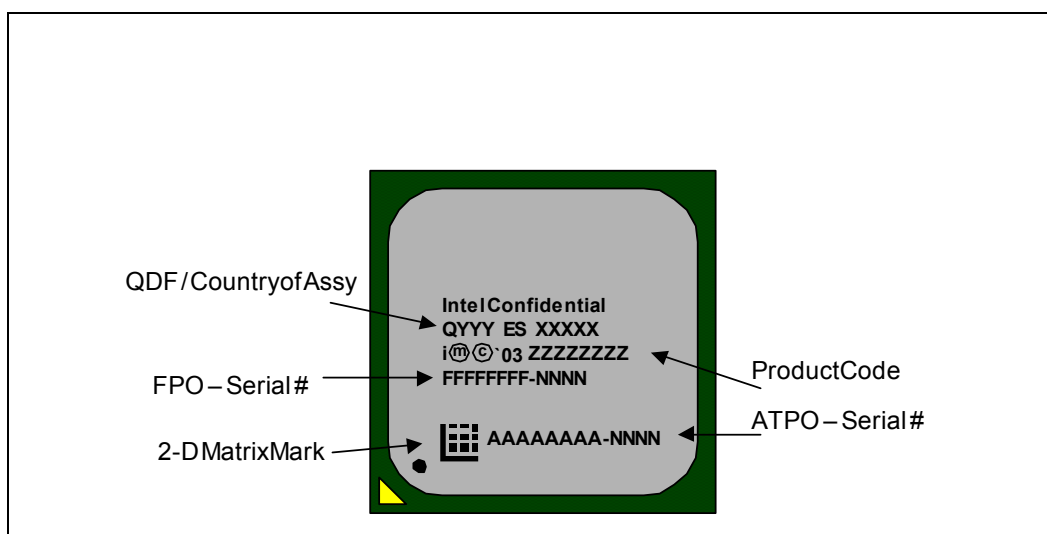



Figure 2. Example of Sample Package Markings





Component Identification Information

The Intel® Celeron® D processor on 90 nm process and in the 478-pin package processor may be identified by the following values.

Family ¹	Model ²
1111b	0011b

NOTES:

1. The Family corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the CUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
2. The Model corresponds to bits [7:4] of the EDX register after RESET, bits [7:4] of the EAX register after the CUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.

Table 1. Intel® Celeron® D Processor on 90 nm Process and in the 478-Pin Package Processor Identification Information

QDF	S-Spec	Core Stepping	Processor Signature	Processor Number	Core Freq (GHz)	Data Bus Freq (MHz)	L2 Cache Size	Processor Package Revision	Package And Revision	Notes
Q611	SL7C4	C0	0F33h	320	2.40	533	256KB	02	478-pin micro-PGA with 35.0 x 35.0 mm FC-mPGA4 package	1
	SL7C5	C0	0F33h	325	2.53	533	256KB	02	478-pin micro-PGA with 35.0 x 35.0 mm FC-mPGA4 package	1
	SL7C6	C0	0F33h	330	2.66	533	256KB	02	478-pin micro-PGA with 35.0 x 35.0 mm FC-mPGA4 package	1
	SL7C7	C0	0F33h	335	2.80	533	256KB	02	478-pin micro-PGA with 35.0 x 35.0 mm FC-mPGA4 package	1

NOTES:

1. These are qualification samples.

Summary Tables of Changes

The following table indicates the Specification Changes, Errata, Specification Clarifications or Documentation Changes that apply to the listed component steppings. Intel intends to fix some of the errata in a future stepping of the component, and to account for the other outstanding issues through documentation or Specification Changes as noted. This table uses the following notations:

Codes Used in Summary Table

Stepping

X:	Erratum, Specification Change or Clarification that applies to this stepping.
(No mark) or (Blank Box):	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.

Status

Doc:	Document change or update that will be implemented.
PlanFix:	This erratum may be fixed in a future stepping of the product.
Fixed:	This erratum has been previously fixed.
NoFix:	There are no plans to fix this erratum.
Shaded:	This item is either new or modified from the previous version of the document.

Note: Each Specification Update item is prefixed with a capital letter to distinguish the product. The key below details the letters that are used in Intel's microprocessor Specification Updates:

A = Intel® Pentium® II processor
 B = Mobile Intel® Pentium® II processor
 C = Intel® Celeron® processor
 D = Intel® Pentium® II Xeon™ processor
 E = Intel® Pentium® III processor
 G = Intel® Pentium® III Xeon™ processor
 H = Mobile Intel® Celeron® processor at 466/433/400/366/333/300 and 266 MHz
 K = Mobile Intel® Pentium® III processor
 L = Unannounced Intel® Celeron® D Processor on 90 nm Process and in the 478-Pin Package
 M = Mobile Intel® Celeron® processor
 N = Intel® Pentium® 4 processor
 O = Intel® Xeon™ processor MP
 P = Intel® Xeon™ processor
 Q = Unannounced Mobile Prescott Processor
 R = Unannounced Intel® Pentium® 4 processor (Prescott)
 S = Unannounced Intel® Xeon™ processor (Nocona)
 T = Mobile Intel® Pentium® 4 Processor-M
 U = Unannounced Intel® Xeon™ processor MP (for future Potomac)
 V = Desktop and Mobile Intel® Celeron® processor on 0.13 Micron Process in Micro-FCPGA Package
 W = Intel® Celeron® M processor (Banias)
 Y = Intel® Pentium® M processor
 Z = Mobile Intel® Pentium® 4 processor with 533 MHz system bus

NO.	B1	C0	D0	Plans	ERRATA
L1	X	X	X	No Fix	Transaction is not retried after BINIT#
L2	X	X	X	No Fix	Invalid opcode 0FFFH requires a ModRM byte
L3	X	X	X	No Fix	Processor may hang due to Speculative Page Walks to Non-Existent System Memory
L4	X	X	X	No Fix	Memory type of the load lock different from its corresponding store unlock
L5	X	X	X	No Fix	Machine check architecture error reporting and recovery may not work as expected
L6	X	X	X	No Fix	Debug mechanisms may not function as expected
L7	X	X	X	No Fix	Cascading of performance counters does not work correctly when forced overflow is enabled
L8	X	X	X	No Fix	EMON event counting of X87 loads may not work as expected
L9	X	X	X	No Fix	System bus interrupt messages without data and which receive a HardFailure response may hang the processor
L10	X	X	X	No Fix	Processor flags #PF instead of #AC on an unlocked CMPXCHG8B instruction
L11	X	X	X	No Fix	FSW may not be completely restored after page fault on FRSTOR or FLDENV instructions

NO.	B1	C0	D0	Plans	ERRATA
L12	X	X	X	No Fix	Processor Issues Inconsistent Transaction Size Attributes for Locked Operation
L13	X	X	X	No Fix	When the processor is in the System Management Mode (SMM), Debug Registers may be fully writeable
L14	X	X	X	No Fix	Shutdown and IERR# may result due to a Machine Check Exception on a Hyper-Threading technology enabled processor
L15	X	X	X	No Fix	Processor may hang under certain frequencies and 12.5% STPCLK# duty cycle
L16	X	X	X	No Fix	System may hang if a fatal cache error causes Bus Write Line (BWL) transaction to occur to the same cache line address as an outstanding Bus Read Line (BRL) or Bus Read-Invalidate Line (BRIL)
L17	X	X	X	No Fix	A write to APIC Task Priority Register (TPR) that lowers priority may seem to have not occurred
L18	X	X	X	No Fix	ITP cannot continue Single Step Execution after the first breakpoint
L19	X	X	X	No Fix	Parity Error in the L1 Cache may cause the processor to hang
L20	X	X		Fixed	BPM4# Signal Not Being Asserted According to Specification
L21	X	X	X	PlanFix	Sequence of locked operations can cause two threads to receive stale data and cause application hang
L22	X			Fixed	Combination of a Bus Lock Target Address and Code in the Same Cache Line May Cause a System Hang
L23	X	X	X	No Fix	A 16-bit Address Wrap Resulting from a Near Branch (Jump or Call) May Cause an Incorrect Address to Be Reported to the #GP Exception Handler
L24	X	X	X	No Fix	Bus Locks and SMC Detection May Cause the Processor to Temporarily Hang
L25	X			Fixed	Signal Reference Voltage for INIT# and STPCLK# is Incorrect
L26	X	X		Fixed	PWRGOOD and TAP Signals Maximum Input Hysteresis Higher Than Specified
L27	X	X	X	Plan Fix	Some Front Side Bus I/O Specifications are not Met
L28	X	X	X	Plan Fix	Incorrect Physical Address Size Returned by CPUID Instruction
L29	X	X	X	No Fix	Incorrect Debug Exception (#DB) May Occur When a Data Breakpoint is set on an FP Instruction
L30	X	X	X	No Fix	xAPIC May Not Report Some Illegal Vector Errors
L31	X	X	X	No Fix	Memory Aliasing of Pages as Uncacheable Memory Type and Write Back (WB) May Hang the System
L32	X	X	X	PlanFix	Interactions Between the Instruction Translation Lookaside Buffer (ITLB) and the Instruction Streaming Buffer May Cause Unpredictable Software Behavior

No.	Plan	Specification Changes
L1	Doc	Privilege Level Change for Monitor/Mwait Instructions

No.	Plan	Specification Clarifications
		There are no Specification Clarifications in this Specification Update revision.

NO.	Plans	DOCUMENTATION CHANGES
		There are no Documentation Changes in this Specification Update revision.

§



Errata

L1. Transaction Is Not Retired after BINIT#

Problem: If the first transaction of a locked sequence receives a HITM# and DEFER# during the snoop phase it should be retried and the locked sequence restarted. However, if BINIT# is also asserted during this transaction, the transaction will not be retried.

Implication: When this erratum occurs, locked transactions will not be retried.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L2. Invalid Opcode 0FFFh Requires A ModRM Byte

Problem: Some invalid opcodes require a ModRM byte and other following bytes, while others do not. The invalid opcode 0FFFh did not require a ModRM in previous generation microprocessors such as Intel® Pentium® II or Pentium III processors, but it is required in the Intel® Pentium® 4 processors, and the Prescott processor.

Implication: The use of an invalid opcode 0FFFh without the ModRM byte may result in a page or limit fault on the Prescott processor. When this erratum occurs, locked transactions will not be retried.

Workaround: To avoid this erratum use ModRM byte with invalid 0FFFh opcode.

Status: For the steppings affected, see the Summary Table of Changes.

L3. Processor May Hang Due to Speculative Page Walks to NonExistent System Memory

Problem: A load operation issued speculatively by the processor that misses the Data Translation Lookaside Buffer (DTLB) results in a page walk. A branch instruction older than the load retires so that this load operation is now in the mispredicted branch path. Due to an internal boundary condition, in some instances the load is not canceled before the page walk is issued.

The Page Miss Handler (PMH) starts a speculative page-walk for the Load and issues a cacheable load of the Page Directory Entry (PDE). This PDE load returns data that points to a page table entry in uncacheable (UC) memory. The PMH issues the PTE Load to UC space, which is issued on the Front Side Bus. No response comes back for this load PTE operation since the address is pointing to system memory, which does not exist.

This load to non-existent system memory causes the processor to hang because other bus requests are queued up behind this UC PTE load, which never gets a response. If the load was accessing valid system memory, the speculative page-walk would successfully complete and the processor would continue to make forward progress.

Implication: Processor may hang due to speculative page walks to non-existent system memory.

Workaround: Page directories and page tables in UC memory space must point to system memory that exists.

Status: For the steppings affected, see the Summary Table of Changes.

L4. Memory Type of the Load Lock Different from its Corresponding Store Unlock

Problem: A use-once protocol is employed to ensure that the processor in a multi-agent system may access data that is loaded into its cache on a Read-for-Ownership operation at least once before it is snooped out by another agent. This protocol is necessary to avoid a multi-agent livelock scenario in which the processor cannot gain ownership of a line and modify it before that data is snooped out by another agent. In the case of this erratum, split load lock instructions incorrectly trigger the use-once protocol. A load lock operation accesses data that splits across a page boundary with both pages of WB memory type. The use-once protocol activates and the memory type for the split halves get forced to UC. Since use-once does not apply to stores, the store unlock instructions go out as WB memory type. The full sequence on the bus is: locked partial read (UC), partial read (UC), partial write (WB), locked partial write (WB). The use-once protocol should not be applied to load locks.

Implication: When this erratum occurs, the memory type of the load lock will be different than the memory type of the store unlock operation. This behavior (Load Locks and Store Unlocks having different memory types) does not however introduce any functional failures such as system hangs or memory corruption.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L5. Machine Check Architecture Error Reporting and Recovery May Not Work as Expected

Problem: When the processor detects errors it should attempt to report and/or recover from the error. In the situations described below, the processor does not report and/or recover from the error(s) as intended.

When a transaction is deferred during the snoop phase and subsequently receives a Hard Failure response, the transaction should be removed from the bus queue so that the processor may proceed. Instead, the transaction is not properly removed from the bus queue, the bus queue is blocked, and the processor will hang.

When a hardware prefetch results in an uncorrectable tag error in the L2 cache, MC0_STATUS.UNCOR and MC0_STATUS.PCC are set but no Machine Check Exception (MCE) is signaled. No data loss or corruption occurs because the data being prefetched has not been used. If the data location with the uncorrectable tag error is subsequently accessed, an MCE will occur. However, upon this MCE, or any other subsequent MCE, the information for that error will not be logged because MC0_STATUS.UNCOR has already been set and the MCA status registers will not contain information about the error which caused the MCE assertion but instead will contain information about the prefetch error event.

When the reporting of errors is disabled for Machine Check Architecture (MCA) Bank 2 by setting all MC2_CTL register bits to 0, uncorrectable errors should be logged in the IA32_MC2_STATUS register but no machine-check exception should be generated. Uncorrectable loads on bank 2, which would normally be logged in the IA32_MC2_STATUS register, are not logged.

When one half of a 64 byte instruction fetch from the L2 cache has an uncorrectable error and the other 32 byte half of the same fetch from the L2 cache has a correctable error, the processor will attempt to correct the correctable error but cannot proceed due to the uncorrectable error. When this occurs the processor will hang.

When an L1 cache parity error occurs, the cache controller logic should write the physical address of the data memory location that produced that error into the IA32_MC1_ADDR REGISTER (MC1_ADDR). In some instances of a parity error on a load operation that hits the L1 cache, however, the cache controller logic may write the physical address from a subsequent load or store operation into the IA32_MC1_ADDR register.

~~The local xAPIC has an Error Status Register, which records all errors it detects. Bit 6 of this register, the Receive Illegal Vector bit, is set when the local xAPIC detects an illegal vector in a message that it received. When an illegal vector error is received on the same internal clock that the error status register is being written due to a previous error, bit 6 does not get set and illegal vector errors are not flagged.~~

When an error exists in the tag field of a cache line such that a request for ownership (RFO) issued by the processor hits multiple tag fields in the L2 cache (the correct tag and the tag with the error) and the accessed data also has a correctable error, the processor will correctly log the multiple tag match error but will hang when attempting to execute the machine check exception handler.

If a memory access receives a machine check error on both 64 byte halves of a 128-byte L2 cache sector, the IA32_MC0_STATUS register records this event as multiple errors, i.e., the valid error

bit and the overflow error bit are both set indicating that a machine check error occurred while the results of a previous error were in the error-reporting bank. The IA32_MC1_STATUS register should also record this event as multiple errors but instead records this event as only one correctable error.

The overflow bit should be set to indicate when more than one error has occurred. The overflow bit being set indicates that more than one error has occurred. Because of this erratum, if any further errors occur, the MCA overflow bit will not be updated, thereby incorrectly indicating only one error has been received.

If an I/O instruction (IN, INS, REP INS, OUT, OUTS, or REP OUTS) is being executed, and if the data for this instruction becomes corrupted, the processor will signal a Machine Check Exception (MCE). If the instruction is directed at a device that is powered down, the processor may also receive an assertion of SMI#. Since MCEs have higher priority, the processor will call the MCE handler, and the SMI# assertion will remain pending. However, while attempting to execute the first instruction of the MCE handler, the SMI# will be recognized and the processor will attempt to execute the SMM handler. If the SMM handler is successfully completed, it will attempt to restart the I/O instruction, but will not have the correct machine state due to the call to the MCE handler. This can lead to failure of the restart and shutdown of the processor.

If PWRGOOD is de-asserted during a RESET# assertion causing internal glitches, the MCA registers may latch invalid information.

If RESET# is asserted, then de-asserted, and reasserted, before the processor has cleared the MCA registers, then the information in the MCA registers may not be reliable, regardless of the state or state transitions of PWRGOOD.

If MCERR# is asserted by one processor and observed by another processor, the observing processor does not log the assertion of MCERR#. The Machine Check Exception (MCE) handler called upon assertion of MCERR# will not have any way to determine the cause of the MCE.

The Overflow Error bit (bit 62) in the IA32_MC0_STATUS register indicates, when set, that a machine check error occurred while the results of a previous error were still in the error reporting bank (i.e. The Valid bit was set when the new error occurred). If an uncorrectable error is logged in the error-reporting bank and another error occurs, the overflow bit will not be set.

The MCA Error Code field of the IA32_MC0_STATUS register gets written by a different mechanism than the rest of the register. For uncorrectable errors, the other fields in the IA32_MC0_STATUS register are only updated by the first error. Any further errors that are detected will update the MCA Error Code field without updating the rest of the register, thereby leaving the IA32_MC0_STATUS register with stale information.

When a speculative load operation hits the L2 cache and receives a correctable error, the IA32_MC1_Status Register may be updated with incorrect information. The IA32_MC1_Status Register should not be updated for speculative loads.

The processor should only log the address for L1 parity errors in the IA32_MC1_Status register if a valid address is available. If a valid address is not available, the Address Valid bit in the IA32_MC1_Status register should not be set. In instances where an L1 parity error occurs and the address is not available because the linear to physical address translation is not complete or an internal resource conflict has occurred, the Address Valid bit is incorrectly set.

The processor may hang when an instruction code fetch receives a hard failure response from the Front Side Bus. This occurs because the bus control logic does not return data to the core, leaving the processor empty. IA32_MC0_STATUS MSR does indicate that a hard fail response occurred.

The processor may hang when the following events occur and the machine check exception is enabled, CR4.MCE=1. A processor that has its STPCLK# pin asserted will internally enter the Stop Grant State and finally issue a Stop Grant Acknowledge special cycle to the bus. If an uncorrectable error is generated during the Stop Grant process it is possible for the Stop Grant special cycle to be issued to the bus before the processor vectors to the machine check handler. Once the chipset receives its last Stop Grant special cycle it is allowed to ignore any bus activity from the processors. As a result, processor accesses to the machine check handler may not be acknowledged, resulting in a processor hang.

Implication: The processor is unable to correctly report and/or recover from certain errors.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L6. Debug Mechanisms May Not Function As Expected

Problem: If the first transaction of a locked sequence receives a HITM# and DEFER# during the snoop phase it should be retried and the locked sequence restarted. However, if BINIT# is also asserted during this transaction, the transaction will not be Certain debug mechanisms may not function as expected on the processor. The cases are as follows:

When the following conditions occur: 1) An FLD instruction signals a stack overflow or underflow, 2) the FLD instruction splits a page-boundary or a 64 byte cache line boundary, 3) the instruction matches a Debug Register on the high page or cache line respectively, and 4) the FLD has a stack fault and a memory fault on a split access, the processor will only signal the stack fault and the debug exception will not be taken.

When a data breakpoint is set on the ninth and/or tenth byte(s) of a floating point store using the Extended Real data type, and an unmasked floating point exception occurs on the store, the break point will not be captured.

When any instruction has multiple debug register matches, and any one of those debug registers is enabled in DR7, all of the matches should be reported in DR6 when the processor goes to the debug handler. This is not true during a REP instruction. As an example, during execution of a REP MOVSW instruction the first iteration a load matches DR0 and DR2 and sets DR6 as FFFF0FF5h. On a subsequent iteration of the instruction, a load matches only DR0. The DR6 register is expected to still contain FFFF0FF5h, but the processor will update DR6 to FFFF0FF1h.

A Data breakpoint that is set on a load to uncacheable memory may be ignored due to an internal segment register access conflict. In this case the system will continue to execute instructions, bypassing the intended breakpoint. Avoiding having instructions that access segment descriptor registers e.g. LGDT, LIDT close to the UC load, and avoiding serialized instructions before the UC load will reduce the occurrence of this erratum.

Implication: Certain debug mechanisms do not function as expected on the processor.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L7. Cascading Of Performance Counters Does Not Work Correctly When Forced Overflow is Enabled

Problem: The performance counters are organized into pairs. When the CASCADE bit of the Counter Configuration Control Register (CCCR) is set, a counter that overflows will continue to count in the other counter of the pair. The FORCE_OVF bit forces the counters to overflow on every non-zero increment. When the FORCE_OVF bit is set, the counter overflow bit will be set but the counter no longer cascades.

Implication: The performance counters do not cascade when the FORCE_OVF bit is set.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L8. EMON Event Counting of x87 Loads May Not Work as Expected

Problem: If a performance counter is set to count x87 loads and floating-point exceptions are unmasked, the FPU Operand (Data) Pointer (FDP) may become corrupted.

Implication: When this erratum occurs, FPU Operand (Data) Pointer (FDP) may become corrupted.

Workaround: This erratum will not occur with floating point exceptions masked. If floating-point exceptions are unmasked, then performance counting of x87 loads should be disabled.

Status: For the steppings affected, see the Summary Table of Changes.

L9. System Bus Interrupt Messages Without Data And Which Receive A HardFailure Response May Hang The Processor

Problem: When a System Bus agent (processor or chipset) issues an interrupt transaction without data onto the System Bus, and the transaction receives a HardFailure response, a potential processor hang can occur. The processor, which generates an inter-processor interrupt (IPI) that receives HardFailure response, will still log the MCA error event cause as HardFailure, even if the APIC causes a hang. Other processors, which are true targets of the IPI, will also hang on hardfailure-without-data, but will not record an MCA HardFailure event as a cause. If a HardFailure response occurs on a System Bus interrupt message with data, the APIC will complete the operation so as not to hang the processor.

Implication: The processor may hang.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L10. Processor Flags #PF Instead of #AC on an Unlocked CMPXC8B Instruction

Problem: If a data page fault (#PF) and alignment check fault (#AC) both occur for an unlocked CMPXC8B instruction, then #PF will be flagged.

Implication: Software that depends #AC before #PF will be affected since #PF is flagged in this case.

Workaround: Remove the software's dependency on the fact that #AC has precedence over #PF. Alternately, if the reload is due to a not present page, reload the page in the page fault handler and then restart the faulting instruction.

Status: For the steppings affected, see the Summary Table of Changes.

L11. FSW May Not Be Completely Restored after Page Fault on FRSTOR or FLDENV Instructions

Problem: If the FPU operating environment or FPU state (operating environment and register stack) being loaded by an FLDENV or FRSTOR instruction wraps around a 64-Kbyte or 4-Gbyte boundary and a page fault (#PF) or segment limit fault (#GP or #SS) occurs on the instruction near the wrap boundary, the upper byte of the FPU status word (FSW) might not be restored. If the fault handler does not restart program execution at the faulting instruction, stale data may exist in the FSW.

Implication: When this erratum occurs, stale data will exist in the FSW.

Workaround: Ensure that the FPU operating environment and FPU state do not cross 64-Kbyte or 4-Gbyte boundaries. Alternately, ensure that the page fault handler restarts program execution at the faulting instruction after correcting the paging problem.

Status: For the steppings affected, see the Summary Table of Changes.

L12. Processor Issues Inconsistent Transaction Size Attributes for Locked Operation

Problem: When the processor is in the Page Address Extension (PAE) mode and detects the need to set the Access and/or Dirty bits in the page directory or page table entries, the processor sends an 8 byte load lock onto the System Bus. A subsequent 8 byte store unlock is expected, but instead a 4 byte store unlock occurs. Correct data is provided since only the lower bytes change, however external logic monitoring the data transfer may be expecting an 8-byte store unlock.

Implication: No known commercially available chipsets are affected by this erratum.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L13. When the Processor Is in the System Management Mode (SMM), Debug Registers May Be Fully Writeable

Problem: When in System Management Mode (SMM), the processor executes code and stores data in the SMRAM space. When the processor is in this mode and writes are made to DR6 and DR7, the processor should block writes to the reserved bit locations. Due to this erratum, the processor may not block these writes. This may result in invalid data in the reserved bit locations.

Implication: Reserved bit locations within DR6 and DR7 may become invalid.

Workaround: Software may perform a read/modify/write when writing to DR6 and DR7 to ensure that the value in the reserved bits are maintained.

Status: For the steppings affected, see the Summary Table of Changes.

L14. Shutdown and IERR# May Result Due to a Machine Check Exception on a Hyper-Threading Technology Enabled Processor

Problem: When a Machine Check Exception (MCE) occurs due to an internal error, both logical processors on a Hyper Threading technology enabled processor normally vector to the MCE handler. However, if one of the logical processors is in the “Wait for SIPI” state, that logical processor will not have a MCE handler and will shut down and assert IERR#.

Implication: A processor with a logical processor in the “Wait for SIPI” state will shut down when an MCE occurs on the other thread.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L15. Processor May Hang Under Certain Frequencies and 12.5% STPCLK# Duty Cycle

Problem: If a system de-asserts STPCLK# at a 12.5% duty cycle, and the processor is running below 2 GHz, and the processor thermal control circuit (TCC) on-demand clock modulation is active, the processor may hang. This erratum does not occur under the automatic mode of the TCC.

Implication: When this erratum occurs, the processor will hang.

Workaround: If use of the on-demand mode of the processor's TCC is desired in conjunction with STPCLK# modulation, then assure that STPCLK# is not asserted at a 12.5% duty cycle.

Status: For the steppings affected, see the Summary Table of Changes.

L16. System May Hang If a Fatal Cache Error Causes Bus Write Line (BWL) Transaction to Occur to the Same Cache Line Address as an Outstanding Bus Read Line (BRL) or Bus Read-Invalidate Line (BRIL)

Problem: A processor internal cache fatal data ECC error may cause the processor to issue a BWL transaction to the same cache line address as an outstanding BRL or BRIL. As it is not typical behavior for a single processor to have a BWL and a BRL/BRIL concurrently outstanding to the same address, this may represent an unexpected scenario to system logic within the chipset.

Implication: The processor may not be able to fully execute the machine check handler in response to the fatal cache error if system logic does not ensure forward progress on the System Bus under this scenario.

Workaround: System logic should ensure completion of the outstanding transactions. Note that during recovery from a fatal data ECC error, memory image coherency of the BWL with respect to BRL/BRIL transactions is not important. Forward progress is the primary requirement.

Status: For the steppings affected, see the Summary Table of Changes.

L17. A Write to APIC Task Priority Register (TPR) That Lowers Priority May Seem to Have Not Occurred

Problem: Uncacheable stores to the APIC space are handled in a non-synchronous way with respect to the speed at which instructions are retired. If an instruction that masks the interrupt flag e.g. CLI is executed soon after an uncacheable write to the TPR that lowers the APIC priority the interrupt masking operation may take effect before the actual priority has been lowered. This may cause interrupts whose priority is lower than the initial TPR but higher than the final TPR to not be serviced until the interrupt flag is finally cleared e.g. STI. Interrupts will remain pending and are not lost.

Implication: This condition may allow interrupts to be accepted by the processor but may delay their service.

Workaround: This can be avoided by issuing a TPR Read after a TPR Write that lowers the TPR value. This will force the store to the APIC priority resolution logic before any subsequent instructions are executed. No commercial operating system is known to be impacted by this erratum.

Status: For the steppings affected, see the Summary Table of Changes.

L18. ITP Cannot Continue Single Step Execution after the First Breakpoint

Problem: ITP will not continue in single step execution after the first software breakpoint. ITP is unable to reset the Resume Flag (RF) bit in the EFLAGS Register.

Implication: The processor repeatedly breaks at the instruction breakpoint address instead of single stepping.

Workaround: Execution after the break will continue, if you manually clear DR7 bit 1 (Global Breakpoint Enable).

Status: For the steppings affected, see the Summary Table of Changes.

L19. Parity Error in the L1 Cache May Cause the Processor to Hang

Problem: If a locked operation accesses a line in the L1 cache that has a parity error, it is possible that the processor may hang while trying to evict the line.

Implication: If this erratum occurs, it may result in a system hang. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L20. BPM4# Signal Not Being Asserted According to Specification.

Problem: BPM4# signal is not being asserted according to the specification. This may cause incorrect operation of In-Target Debuggers particularly at higher FSB frequencies.

Implication: In-Target Debuggers may not function at higher than 133/533 MHz FSB.

Workaround: One method is to reduce the FSB common clock frequency to 133 MHz or lower. For higher FSB speeds, In-Target Debuggers have a built-in function (test2010) that tells the hardware to ignore BPM4# assertions. This may degrade the debugger performance but will give correct results.

Status: For the steppings affected, see the Summary Table of Changes.

L21. Sequence of Locked Operations Can Cause Two Threads to Receive Stale Data and Cause Application Hang

Problem: While going through a sequence of locked operations, it is possible for the two threads to receive stale data. This is a violation of expected memory ordering rules and causes the application to hang.

Implication: When this erratum occurs in an HT Technology enabled system, an application may hang.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the Summary Table of Changes.

L22. Combination of a Bus Lock Target Address and Code in the Same Cache Line May Cause a System Hang

Problem: If bus lock target addresses and code are combined in the same cache line, it is possible that a continuous stream of bus locks from one thread (attempting to acquire or release the lock) may prevent the execution of code on the second thread (that is also attempting to release or acquire the same lock.)

Implication: When this erratum occurs in an HT Technology enabled system, it may cause a system hang.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L23. A 16-bit Address Wrap Resulting from a Near Branch (Jump or Call) May Cause an Incorrect Address to Be Reported to the #GP Exception Handler

Problem: If a 16-bit application executes a branch instruction that causes an address wrap to a target address outside of the code segment, the address of the branch instruction should be provided to the general protection exception handler. It is possible that, as a result of this erratum, that the general protection handler may be called with the address of the branch target.

Implication: A 16-bit software environment that is affected by this erratum will see that the address reported by the exception handler points to the target of the branch rather than the address of the branch instruction.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes.

L24. Bus Locks and SMC Detection May Cause the Processor to Temporarily Hang

Problem: The processor may temporarily hang in an HT Technology enabled system, if one logical processor executes a synchronization loop that includes one or more bus locks and is waiting for release by the other logical processor. If the releasing logical processor is executing instructions that are within the detection range of the self modifying code (SMC) logic, then the processor may be locked in the synchronization loop until the arrival of an interrupt or other event.

Implication: If this erratum occurs in an HT Technology enabled system, the application may temporarily stop making forward progress. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes at the beginning of this section.

L25. Signal Reference Voltage for INIT# and STPCLK# Is Incorrect

Problem: The input levels for the Asynchronous GTL+ Signal Group DC are specified to be derived from $V_{cc}/2$. The INIT# and STPCLK# input levels are instead derived from GTLREF.

Implication: The INIT# and STPCLK# inputs may switch at different levels than previously documented specifications. Intel has not observed any issues in validation or simulation as a result of this erratum.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes at the beginning of this section.

L26. PWRGOOD and TAP Signals Maximum Input Hysteresis Higher Than Specified

Problem: The maximum input hysteresis for the PWRGOOD and TAP input signals is specified at 350 mV. The actual value could be as high as 800 mV.

Implication: The PWRGOOD and TAP inputs may switch at different levels than previously documented specifications. Intel has not observed any issues in validation or simulation as a result of this erratum.

Workaround: None identified.

Status: For the steppings affected, see the Summary Table of Changes at the beginning of this section.

L27. Some Front Side Bus I/O Specifications Are Not Met

Problem: The following front side bus I/O specifications are not met:

- The $V_{IH(min)}$ for the GTL+ signals is specified as $GTLREF + (0.10 * V_{CC})$ [V].
- The $V_{IH(min)}$ for the Asynchronous GTL+ signals is specified as $V_{cc}/2 + (0.10 * V_{CC})$ [V].
- Common Clock Output Valid Delay (min) is specified as -0.250 ns.
- Common Clock Input Setup Time is specified as 0.700 ns.
- Source Synchronous Input Setup Time to Strobe is specified as 0.150 ns.
- Source Synchronous Input Hold Time to Strobe is specified as 0.150 ns.

Implication: This erratum can cause functional failures depending upon system bus activity. It can manifest itself as data parity, address parity, and/or machine check errors.

Workaround: Due to this erratum, the system should meet the following voltage levels and processor timings:

- The $V_{IH(min)}$ for GTL+ signals is now $GTLREF + (0.20 * V_{CC})$ [V].
- The $V_{IH(min)}$ for the Asynchronous GTL+ signals is now $V_{cc}/2 + (0.20 * V_{CC})$ [V].
- Common Clock Output Valid Delay (min) is now -0.300 ns.
- Common Clock Input Setup Time is now 0.900 ns.
- Source Synchronous Input Setup Time to Strobe is now 0.350 ns.
- Source Synchronous Input Hold Time to Strobe is now 0.350 ns.

Status: For the steppings affected, see the *Summary Table of Changes*.

L28. Incorrect Physical Address Size Returned by CPUID Instruction

Problem: The CPUID instruction Function 80000008H (Extended Address Sizes Function) returns the address sizes supported by the processor in the EAX register. This Function returns an incorrect physical address size value of 40 bits. The correct physical address size is 36 bits.

Implication: Function 80000008H returns an incorrect physical address size value of 40 bits.

Workaround: None identified.

Status: For the steppings affected, see the Summary *Table of Changes*.

L29. Incorrect Debug Exception (#DB) May Occur When a Data Breakpoint is set on an FP Instruction

Problem: The default Microcode Floating Point Event Handler routine executes a series of loads to obtain data about the FP instruction that is causing the FP event. If a data breakpoint is set on the instruction causing the FP event, the load in the microcode routine will trigger the data breakpoint resulting in a Debug Exception.

Implication: An incorrect Debug Exception (#DB) may occur if data breakpoint is placed on an FP instruction. Intel has not observed this erratum with any commercially available software or system.

Workaround: None identified.

Status: For the steppings affected, see the Summary *Table of Changes*.

L30. xAPIC May Not Report Some Illegal Vector Errors

Problem: The local xAPIC has an Error Status Register, which records all errors. The bit 6 (the Receive Illegal Vector bit) of this register, is set when the local xAPIC detects an illegal vector in a received message. When an illegal vector error is received on the same internal clock that the error status register is being written (due to a previous error), bit 6 does not get set and illegal vector errors are not flagged.

Implication: The xAPIC may not report some Illegal Vector errors when they occur at approximately the same time as other xAPIC errors. The other xAPIC errors will continue to be reported.

Workaround: None identified.

Status: For the steppings affected, see the Summary *Table of Changes*.

L31. Memory Aliasing of Pages as Uncacheable Memory Type and Write Back (WB) May Hang the System

Problem: When a page is being accessed as either Uncacheable (UC) or Write Combining (WC) and WB, under certain bus and memory timing conditions, the system may loop in a continual sequence of UC fetch, implicit writeback, and Request for Ownership (RFO) retries.

Implication: This erratum has not been observed in any commercially available operating system or application. The aliasing of memory regions, a condition necessary for this erratum to occur, is documented as being unsupported in the *IA-32 Intel® Architecture Software Developer's Manual*, Volume 3, section 10.12.4, Programming the PAT. However, if this erratum occurs the system may hang.

Workaround: The pages should not be mapped as either UC or WC and WB at the same time.

Status: For the steppings affected, see the Summary *Table of Changes*.

L32. Interactions between the Instruction Translation Lookaside Buffer (ITLB) and the Instruction Streaming Buffer May Cause Unpredictable Software Behavior

Problem: Complex interactions within the instruction fetch / decode unit may make it possible for the processor to execute instructions from an internal streaming buffer containing stale or incorrect information.

Implication: When this erratum occurs, an incorrect instruction stream may be executed resulting in unpredictable software behavior.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the Summary *Table of Changes*.

§

Specification Changes

The Specification changes listed in this section apply to the following documents:

- Prescott New Instruction Software Developer's Guide

All Specification Changes will be incorporated into a future version of the appropriate Intel® Celeron® D Processor on 90 nm Process and in the 478-Pin Package Processor documentation.

L1. Privilege Level Change for the Monitor/Mwait Instructions

The Monitor and Mwait instructions were originally defined in the Prescott New Instruction Software Developer's Guide to be unconditionally executable from all privilege levels, 0 through 3. This specification has been changed so that processors having the Monitor/Mwait feature flag set in CPUID allow the execution of the Monitor/Mwait instructions at privilege levels 1 through 3, with unconditional execution supported at privilege level 0. Application software should test for the appropriate support for these instructions by attempting to execute them prior to their unconditional execution. If the instructions are not supported an invalid opcode exception will be generated.

§



Specification Clarifications

There are no specification clarifications in this Specification Update revision.

§



Documentation Changes

There are no documentation changes in this Specification Update revision.

§